

**METHOD, SYSTEM, AND COMPUTER PROGRAM PRODUCT FOR
COMPUTER SYSTEM VULNERABILITY ANALYSIS AND FORTIFICATION**

Background of the Invention

Field of the Invention

[0001] This invention relates to computer security and, more particularly, to vulnerability analysis systems.

Description of the Related Art

[0002] Many of today's software solutions (programs) are complex applications composed of a large number of components interacting together to solve customer problems. For example, an enterprise software solution may be composed of a number of large components, such as DB2, IBM HTTP Server, WebSphere Application Server, and WebSphere Portal Server. In such a solution, most of the major components are full-scale applications in their own right and cross-component security administration (e.g., the process of making sure all of the entry-points and components of a business system meet mandatory security requirements) is a complex, time-consuming task. For example, applying the latest patches, having secure authentication methods for login, encryption for data communication, etc., especially in environments consisting of multiple interacting platforms, operating

systems, and products enforcing the security policies, can be very time consuming and error-prone. New security vulnerabilities are often discovered in the individual components and fixes to them are released by the vendors, but there is no solution for automatic application of the patches to the individual components and to the software solution as a whole.

[0003] Current software solutions require manual auditing to ensure that no serious security issues exist. Some security applications (vulnerability analysis systems) do exist that scan software solutions and the systems on which they run for weaknesses (e.g., open TCP/IP ports, etc.) that are typically found in all software solutions/systems. They use generic techniques such as port scanning, packet flooding, and the like to perform what is known as “ethical hacking,” that is, attempting to gain access to a system using known vulnerabilities, to see if such vulnerabilities are present in the system under analysis. One example of such a system is SATAN (Security Analysis Tool for Auditing Networks), a free software package released in the 1990's.

[0004] With software solutions increasing in size and complexity, it is increasingly difficult to keep up with the task of identifying and fixing security issues manually. These “virtual hackers”, while identifying some system weaknesses, are not “creative” in that they

rely on generic, known attack methods (such as the previously-mentioned port scanning method) that are based not on the software solutions being analyzed but instead on attacks that have been shown to work on many systems in the past. Further, these systems do not have mechanisms to fix the identified security issues automatically. For example, intrusion detection systems of the prior art detect intrusions and then alert the users as to the need to react to the intrusion. Also, in a complex software solution, even if the individual components of the software solution are secure, it cannot be assumed that the combination of these components is also secure.

[0005] Accordingly, it would be desirable to have a system to detect and fix security holes automatically in software solutions before an intrusion occurs, and that can identify vulnerabilities beyond the generic vulnerabilities found by prior art systems, and that can correlate security risks across multiple components in a software solution.

Summary of the Invention

[0006] The present invention is a method, system, and computer program product for the automatic detection and fixing of security vulnerabilities in both individual software components and across complex, multi-component software solutions. In accordance with

the present invention, the architecture of the software solution to be monitored is analyzed prior to its being monitored. The present invention uses data derived from the analysis to proactively identify possible ways to attack the software solution. It then periodically scans the software solution being monitored and the system on which it runs, and attempts attacks on it. A list of possible attacks is continuously updated, for example, in a manner similar to virus signatures provided by virus security companies, and a log is generated describing which attacks were successful and which ones failed. At the end of each scan it can generate a report that identifies the weaknesses found in the solution. If the weaknesses have a known remedy, then the present invention can apply the remedy automatically (according to the set of rules embodied in a “fix” policy). An administrator is also informed and the event is logged.

[0007] In an exemplary embodiment, the present invention comprises a method for detection and correction of security vulnerabilities in a computing environment, comprising: analyzing a software solution to identify legal and illegal external interfaces thereto; attempting to access said software solution using the identified illegal external interfaces; and storing a record of any illegal external interfaces that allow access to said software solution.

Brief Description of the Drawings

[0008] Figure 1 is a block diagram of the general architecture of the present invention in an environment in which it may be used;

[0009] Figure 2 is an example of a directed graph illustrating an example of a portion of a bank transaction involving the electronic transfer of funds.

Detailed Description of the Preferred Embodiments

[0010] Figure 1 is a block diagram of the general architecture of the present invention in an environment in which it may be used. Referring to Figure 1, a vulnerability analysis and fortification tool (VAF) 102 is coupled to an API/attacks/reports repository 103. Repository 103 stores detailed information regarding how to make API calls for the software system(s) under analysis. Repository 103 is central to the operation of VAF 102 and contains XML files (explained in more detail below) that provide all of the information necessary to identify valid (legal) and invalid (illegal) methods for access to a system, and to launch a vulnerability attack on the system and determine when the attack is successful. More particularly, the repository 103 stores information related to what different external API calls

(external interfaces) have been exposed by the system, what are the valid arguments and return codes for such calls, and what constitutes a successful and legal method call. This information may include access sequences (password sequences and return codes, preconditions required for access, etc.) for software solutions being monitored by VAF 102.

[0011] In a preferred embodiment, for each software solution being monitored, a VAF agent is utilized. For example, a VAF agent 104 monitors software solution 112 and VAF agent 106 monitors software solution 114. Each VAF agent has a local database coupled thereto (local database 108 and local database 110) which is used to store the XML descriptions described above, historical data regarding attacks, and attack patterns that have a high success rate. Separate VAF agents are used because the system to be protected could be a distributed system scattered over many locations. In such cases, a VAF agent performs the security analysis of the subcomponent assigned to it and sends the results over to the VAF 102.

[0012] A prior auditing system 116 (e.g., the previously described SATAN system) can be used in conjunction with the VAF 102. In this configuration, the prior art auditing system 116 monitors the system for common, previously-known weaknesses, while the VAF

102 (and its VAF agents) monitors and corrects the system for specific weaknesses identified by the VAF 102.

[0013] To keep the VAF up to date with the most current security information, VAF 102 can be coupled to a security server 126 via an Internet connection 124 and a firewall intrusion detection system 118. In a preferred embodiment, there are separate firewall/intrusion detection systems 120 and 122 provided for each software solution being monitored by VAF 102.

[0014] Using the system illustrated in Figure 1, the VAF 102 obtains detailed information regarding the architecture of each software solution being monitored. The VAF obtains knowledge of the external interfaces to a solution (e.g., all mechanisms exposed by the system to communicate with external applications, e.g., APIs, Webservices, HTML forms, etc.) by using an XML description of its interfaces stored in API repository 103. These XML descriptions can be created specifically for use by the VAF and/or can be migrated from other established sources, e.g., existing test harnesses have essentially the same information, although possibly in a form requiring modification to be “understood” by the VAF.

[0015] For each interface described, there is a set of preconditions that have to be met before “legal” invocation of the interface, and also a corresponding acceptable return code or codes indicating success or failure to access the interface. Based on the XML description of the interfaces to the software solutions, the VAF creates a directed graph, with the nodes being a state in the system and the edges being method invocations. The graph is simply a graphical representation, somewhat like a map, showing how a system may be legally accessed, as well as helping identify potential methods of illegally accessing the system. Any manner of identifying how a system may be legally accessed will suffice; a directed graph is one example of how it can be done.

[0016] An example of such a directed graph is illustrated in Figure 2, illustrating an example of a portion of a bank transaction involving an electronic transfer of funds. Referring to Figure 2, nodes are represented by oval shapes and method invocations are indicated by solid arrows connecting the various nodes. For example, referring to Figure 2, from an initial state 202, a method 212 is invoked to successfully traverse to a first state 204. Method 212 could comprise, for example, a precondition of the input of a valid user ID and

password, and the return of a valid return code indicating the receipt of a valid user ID and password, i.e., method 212 could represent a typical logon process for an ATM machine.

[0017] From the first state 204, a method 214 could be invoked having as a precondition, first, the execution of method 212 and the return of a valid return code, followed by the submission of a second pass phrase (for additional security) and the return of an additional valid return code (e.g., a token) indicating the receipt of a valid pass phrase.

[0018] Once in state 206, multiple options are available. For example, method 216 could be invoked, whereby the user is given the option of changing their password by entering a new password in a new password screen, and the successful return of a valid return code indicating the inputting of a valid new password. Alternatively, or in addition, method 218 can be invoked, whereby the user invokes an immediate request to perform a transfer of funds. Alternatively, the user can first invoke method 220 and obtain an account balance (and enter state 208), followed by a method 222, which is a request to transfer funds.

[0019] Each of these valid or “legal” external interfaces with the banking program are illustrated in Figure 2 by solid arrows connecting to the various nodes. To summarize, the

legal sequence of API calls that are required to transfer money from one account to another are as follows:

- [0020] 1. Login using a valid username and password (transition from initial state to state 1);
- [0021] 2. Once logged in, as another level of security, a secret password phrase must be entered (transition from state 1 to state 2);
- [0022] 3. Once in state 2, the initial password can be reset (stay in state 2); the account balance for a particular account can be obtained (states 2-3); or money can be transferred from one account to another (states 2-4).

[0023] All other transitions are illegal; for example, a transfer from the initial state to state 210 (illustrated by dotted transition line 224), would allow a transfer money between accounts without having logged in, which is an illegal operation. Other illegal transitions are illustrated by dotted transition lines 226, 228, and 230.

[0024] The following is an example of an XML description that describes the API represented by the directed graph of Figure 2:

```
<?xml version="1.0" encoding="UTF-8"?>
<method name="login" id="1">
  <parameters>
    <param name="username">
      neeraj
    </param>
    <param name="password">
```

```
        pwd
        </param>
    </parameters>
    <precondition> </precondition>
    <validreturnvalue>
        true
    </validreturnvalue>
</method>
<method name="pwdphrase" id="2">
    <parameters>
        <param name="passphrase">
            abracadabra
        </param>
    </parameters>
    <precondition>
        <method id="1">
            true
        </method>
    </precondition>
</method>
```

[0025] In the above XML snippet method “pwdphrase” has a precondition that method login with id=1 should have returned a value= true before pwdphrase can be legally executed. This example illustrates how all of the interfaces of a system can be described in XML, along with the dependencies between them.

[0026] The VAF tool of the present invention can also log the presence of known threat patterns along with those that have not yet been identified as security threats. The VAF tool scans the system for vulnerabilities regularly. When it discovers a vulnerability, it does the following. If there is a known remedy (e.g., a “patch”) to fix the security hole,

that remedy is applied. If no known remedy exists, the administrator is alerted of the vulnerability. Next, the pattern of calls that lead to the discovery of the security hole is logged. This logged history is later used to optimize vulnerability detection based on the pattern of API calls that identify the vulnerability. This newly discovered information is also made available to all other VAF systems to download via security servers. If a new security hole is discovered, the pattern of calls required to breach the security is made available to all the VAF systems. These systems then check whether such a pattern of calls is possible on the specific system that their particular VAF is monitoring and, if possible, alert the administrator or apply the patch (if it is available). Thus, in the future, if these patterns are identified as being vulnerable to attacks, the application would know exactly where the fixes need to be applied.

[0027] The VAF of the present invention is superior to other current security solutions, such as an intrusion detection system (IDS), in that it analyzes the software being monitored before monitoring to discover vulnerability, and periodically updates itself with the latest attack techniques and defenses from other sources. It then tries these attacks on systems it is responsible for monitoring. If any attack is successful, and if it has a known solution, VAF automatically applies the fix; otherwise, it alerts the user of its vulnerability,

which helps to enable the development of remedies for attacks. These remedies then will be available in the future for automatic application by the VAF.

[0028] One advantage of the present invention is that it can be packaged as an add-on component that can be shipped along with an enterprise software solution. Having a protective component added to the solution will improve the solutions marketability. VAF can also be used in conjunction with other automatic computing technologies, such as Common Log and Trace and ABLE, or it can function alone. This allows enterprise customers the maximum flexibility when configuring their computer systems.

[0029] The above-described steps can be implemented using standard well-known programming techniques. The novelty of the above-described embodiment lies not in the specific programming techniques but in the use of the steps described to achieve the described results. Software programming code which embodies the present invention is typically stored in permanent storage of some type, such as permanent storage of the VAF itself. In a client/server environment, such software programming code may be stored with storage associated with a server. The software programming code may be embodied on any of a variety of known media for use with a data processing system, such as a diskette, or hard

drive, or CD-ROM. The code may be distributed on such media, or may be distributed to users from the memory or storage of one computer system over a network of some type to other computer systems for use by users of such other systems. The techniques and methods for embodying software program code on physical media and/or distributing software code via networks are well known and will not be further discussed herein.

[0030] It will be understood that each element of the illustrations, and combinations of elements in the illustrations, can be implemented by general and/or special purpose hardware-based systems that perform the specified functions or steps, or by combinations of general and/or special-purpose hardware and computer instructions.

[0031] These program instructions may be provided to a processor to produce a machine, such that the instructions that execute on the processor create means for implementing the functions specified in the illustrations. The computer program instructions may be executed by a processor to cause a series of operational steps to be performed by the processor to produce a computer-implemented process such that the instructions that execute on the processor provide steps for implementing the functions specified in the illustrations. Accordingly, the figures support combinations of means for performing the specified

functions, combinations of steps for performing the specified functions, and program instruction means for performing the specified functions.

[0032] Although the present invention has been described with respect to a specific preferred embodiment thereof, various changes and modifications may be suggested to one skilled in the art and it is intended that the present invention encompass such changes and modifications as fall within the scope of the appended claims.